Inference on filtered and smoothed probabilities in Markov-switching autoregressive models*

Rocio Alvarez
Universidad Central de Chile
rocio.alvareza@ucentral.cl

Maximo Camacho[†]
Universidad de Murcia
mcamacho@um.es

Manuel Ruiz
Universidad Politécnica de Cartagena
manuel.ruiz@upct.es

September 1, 2017

Abstract

We derive a statistical theory that provides useful asymptotic approximations to the distributions of the single inferences of filtered and smoothed probabilities, derived from time series characterized by Markov-switching dynamics. We show that the uncertainty in these probabilities diminishes when the states are separated, the variance of the shocks is low, and the time series or the regimes are persistent. As empirical illustrations of our approach, we analyze the U.S. GDP growth rates and the U.S. real interest rates. For both models, we illustrate the usefulness of the confidence intervals when identifying the business cycle phases and the interest rate regimes.

Keywords: Markov Switching, Business Cycles, Time Series.

JEL Classification: E32, C22, E27.

*We would like to thank M. Carnero, T. Niguez, G. Perez Quiros, P. Poncela, and E. Ruiz, the editor, the associate editor and three anonymous reviewers for their helpful comments and suggestions. M. Camacho and M. Ruiz acknowledge the financial support from projects ECO2016-76178-P and ECO2015-65637-P, respectively. Rocio Alvarez acknowledges the financial support from project CIP16013 (Universidad Central de Chile). This study is the result of the activity carried out under the program Groups of Excellence of the region of Murcia, the Fundación Séneca, Science and Technology Agency of the region of Murcia project 19884/GERM/15. All remaining errors are our responsibility.

[†]Corresponding Author: Universidad de Murcia, Facultad de Economia y Empresa, Departamento de Métodos Cuantítativos para la Economía y la Empresa, 30100, Murcia, Spain. E-mail: mcamacho@um.es.

1 Introduction

In Markov-switching autoregressive models, the behavior of a time series is assumed to be disrupted occasionally by shocks that produce different dynamics, regimes, or states, the transitions of which are governed by an unobservable state variable that is usually assumed to follow a first-order Markov chain. These models have become extremely popular in economic analyses since the work of Hamilton (1989). Among other areas, regime-switching models have been applied to exchange rate swings (Engel and Hamilton, 1990; Engel, 1994), stock market returns (Pagan and Schwert, 1990; Hamilton and Lin, 1996), interest rates (Garcia and Perron, 1996; Dahlquist and Gray, 2000.), asset returns (Kim, Nelson and Startz, 1998; Ang and Bekaert, 2002), and asymmetries over the business cycle (McConnell and Perez-Quiros, 2000; Chauvet and Hamilton, 2006).[1]

The widespread success of these models might rely on the fact that analysts can easily draw a probabilistic inference about the hidden Markov chain, given the observations of the time series. The inference used to separate the unobserved regimes is expressed through the filtered state probabilities and the smoothed full-sample probabilities, and historical dates of the turning points are typically established when the probability of one particular regime crosses some pre-specified thresholds (Chauvet and Hamilton, 2006). This is of interest in its own right, because the regimes can provide substantive meaning about the occurrence of economic events.

Although the outcomes of these analyses readily lead to policy implications, note that they are developed from single values of filtered and smoothed inferences, which do not reflect the uncertainty surrounding these inferences. In this context, the primary purpose of this study is to derive a statistical theory that provides useful asymptotic approximations to the distributions of the single inferences of filtered and smoothed probabilities, derived from time series characterized by Markov-switching dynamics.

Bickel, Ritov, and Ryden (1998) show that under mild conditions, the inverse of the negative of the Hessian evaluated at the Maximum-Likelihood (ML) estimates is a consistent estimator of the covariance matrix of the set of parameters governing the Markov-switching autoregressive models. Using these results, we show how the delta method gives a technique for computing the variance of these transformations based on a Taylor series approximation, which is used to compute their respective confidence intervals.

To evaluate the performance of our proposal, we conduct a set of Monte Carlo simulations, allowing regimes affecting the mean, the autoregressive parameters, and variance. The results suggest a thought-provoking, recurrent pattern: the uncertainty on filtered and smoothed probabilities increases significantly around the turning points. Therefore, we propose using sharp increases in the uncertainty to detect phase changes and to separate the regimes, which could be especially useful in real-time analyses. In addition,

[1] Although we focus on economic applications, the Markov-switching model has found widespread applications in many areas, including bioinformatics, biology, finance, hydrology, marketing, medicine, and speech recognition. See Fruhwirth-Schnatter (2006) for an overview.

we show that the uncertainty diminishes when the regimes are clearly separated (i.e., the within-state means are far from each other), the variance of the idiosyncratic shocks is low, the inertia of the time series is high, and the regimes are persistent.

To illustrate our approach empirically, we propose a twofold exercise. In the first exercise, we compute inferences on the U.S. business cycle phases using a model where GDP is characterized by a recurrent sequence of shifts between two fixed equilibria of high-growth and low-growth means. The in-sample results suggest that establishing the historical turning points when the probability of one particular business cycle state remains above or below some pre-specified thresholds, without knowing the uncertainty surrounding the point estimates, is inadvisable, particularly when the point estimates are close to the thresholds. In a real-time analysis, we find that a salient characteristic of the U.S. cycle dynamics is that the uncertainty of the filtered and smoothed probabilities increases significantly when the changes in the business cycle phase occur, while it decreases significantly in the course of the new phase. Therefore, it is worthwhile considering the confidence intervals on the business cycle probabilities when monitoring ongoing economic developments.

In the second empirical example, we follow Garcia and Perron (1996) and allow three possible regimes affecting both the mean and variance of the U.S. ex-post real interest rate, which follows autoregressive dynamics of order two. Based on an updated sample, our results still suggest that the interest rate is characterized by three distinct phases of low, middle, and high levels. We find single distinct phase changes from the low regime to the middle regime in 1953.3, and from the middle regime to the low regime in 1973.4. However, dating the regime shifts becomes increasingly uncertain after this point. The shift to the high regime could occur between 1980.4 and 1981.2, the shift to the middle regime could occur between 1986.4 and 1990.2, and the shift to the low regime could occur between 2002.1 and 2003.1.

The remainder of this paper is organized as follows. Section 2 presents the asymptotic distribution theory for the filtered and smoothed probabilities of one particular regime, and outlines a simple method to construct asymptotically valid confidence intervals. Section 3 proposes a simulation experiment to assess the impact on the uncertainty surrounding the filtered and smoothed probabilities of the model parameters. Section 4 shows the usefulness of the proposed model by analyzing the U.S. business cycle phases and the U.S. interest rate regimes. Lastly, Section 5 concludes the paper.

2 Computing uncertainty

2.1 The model

Let y_t be a stationary time series of $T + p$ observations whose autoregressive dynamics evolve according to an unobservable K-state Markov-chain process s_t.[2] For the sake of generality, the means, regression

[2] General characterizations of stationarity conditions for such processes can be found in Francq and Zakoïan (2001).

coefficients, and volatility of the Markov-switching autoregressive model are state-dependent:

$$y_t = \mu_{s_t} + \sum_{j=1}^{p} \phi_{j,s_{t-j}} (y_{t-j} - \mu_{s_{t-j}}) + \varepsilon_t, \tag{1}$$

where $\varepsilon_t \sim N(0, \sigma^2_{s_t})$ and p is the lag-length of the underlying state-dependent autoregressive process.[3] To complete the statistical characterization of this process, we assume that s_t is a Markov chain of order one. Then, the probability of a change in regime depends on the past only through the value of the most recent regime

$$P(s_t = j | s_{t-1} = i, ..., s_1 = l, Y_{t-1}) = P(s_t = j | s_{t-1} = i) = p_{ij}, \tag{2}$$

where $Y_t = y_1, y_2, ..., y_t$, and $i, j = 0, 1, \ldots, K-1$.

Since the nonlinear autoregressive process depends not only on s_t, but also on $s_{t-1}, ..., s_{t-p}$, it is convenient to define the latent variable $s^*_t = (s_t, s_{t-1}, \ldots, s_{t-p})$, which results in K^{p+1} different states. The transition probabilities of s^*_t can easily be found from the transition probabilities of the primitive states s_t. Let us define the states $\mathbf{j}$ of s^*_t as $\mathbf{j} = (j_0, j_1, \ldots, j_p)$, with $j_i \in \{0, 1, \ldots, K-1\}$, $i = 0, 1, ..., p$. Then, the transition probabilities of s^*_t are

$$P(s^*_t = \mathbf{j} | s^*_{t-1} = \mathbf{i}) := p^*_{\mathbf{ij}} = \begin{cases} p_{i_0 j_0} & \text{always } i_r = j_{r-1} \text{ for } r = 1, 2 \ldots, p \\ \\ 0 & \text{otherwise} \end{cases} \tag{3}$$

2.2 Variance of filtered probabilities

We collect the r model parameters in the vector $\theta = (\mu_0, \ldots, \ \mu_{K-1}, \phi_{1,0}, \ldots, \phi_{1,K-1}, \phi_{p,0}, \ldots, \phi_{p,K-1}, \sigma^2_0, \ldots, \sigma^2_{K-1}, p_{00}, \ldots, p_{K-1K-1})$. As an application of Bayes' law for this setting, Hamilton (1989) computes estimates of the filtered probabilities, as follows:

$$P(s^*_t = \mathbf{i} | \theta, Y_t) := P^{\mathbf{i}}_t(\theta) = \frac{f(y_t | \theta, Y_{t-1}, s^*_t = \mathbf{i}) P(s^*_t = \mathbf{i} | \theta, Y_{t-1})}{f(y_t | \theta, Y_{t-1})}, \tag{4}$$

where $f(y_t|\theta, Y_{t-1}, s^*_t)$ is the pdf of a Gaussian distribution with mean $\mu_{s_t} + \sum\limits_{j=1}^{p} \phi_{s_{t-j}} (y_{t-j} - \mu_{s_{t-j}})$ and variance $\sigma^2_{s_t}$, and $f(y_t|\theta, Y_{t-1}) = \sum\limits_{\mathbf{i}} f(y_t|\theta, Y_{t-1}, s^*_t = \mathbf{i}) P(s^*_t = \mathbf{i}|\theta, Y_{t-1})$.

Then, given that the random series s_t follows a first-order Markov-chain process, it is easy to see that the following equalities hold:

$$P(s^*_t = \mathbf{i} | \theta, Y_{t-1}) = \sum_{\mathbf{j}} p^*_{\mathbf{ji}} P^{\mathbf{j}}_{t-1}(\theta), \tag{5}$$

and

$$f(y_t | \theta, Y_{t-1}) = \sum_{\mathbf{i}} f(y_t | \theta, Y_{t-1}, s^*_t = \mathbf{i}) \sum_{\mathbf{j}} p^*_{\mathbf{ji}} P^{\mathbf{j}}_{t-1}(\theta). \tag{6}$$

[3]Following the standard assumptions on Markov-switching autoregressive models, we focus on normal errors in this study. However, this is not restrictive, and can easily be generalized.

From (6) and (5), if we denote $\psi_t^{\mathbf{i}}$ as the pdf of a Gaussian distribution with mean $\mu_{i_0} + \sum_{z=1}^{p} \phi_{i_z}(y_{t-z} - \mu_{i_z})$ and variance $\sigma_{i_0}^2$, the filtered probabilities become

$$P_t^{\mathbf{i}}(\theta) = \frac{\psi_t^{\mathbf{i}} \sum_{\mathbf{j}} p_{\mathbf{ji}}^* P_{t-1}^{\mathbf{j}}(\theta)}{\sum_{\mathbf{i}} \psi_t^{\mathbf{i}} \sum_{\mathbf{j}} p_{\mathbf{ji}}^* P_{t-1}^{\mathbf{j}}(\theta)}. \tag{7}$$

Using the results suggested by Hamilton (1989), a natural estimate of the vector of parameters θ is the ML estimator $\widehat{\theta}$, because the conditional likelihood function can be obtained as a by-product of Hamilton's filter.[4] Bickel, Ritov, and Ryden (1998) show that under mild conditions the ML estimator is not only consistent, but is also asymptotically normal, and prove that the observed information matrix is a consistent estimator of the Fisher information. Thus, the distribution of the ML estimate $\widehat{\theta}$ can be well approximated for large samples by

$$\widehat{\theta} \sim N(\theta, T^{-1} I_{\widehat{\theta}}^{-1}), \tag{8}$$

where $I_{\widehat{\theta}}$ is known as the Fisher information matrix. In practice, the covariance matrix can be estimated by

$$I_{\widehat{\theta}} = -T^{-1} \frac{\partial \mathcal{L}(\theta)}{\partial \theta \partial \theta'} |_{\theta = \widehat{\theta}}, \tag{9}$$

where $\mathcal{L}(\theta)$ is the log-likelihood

$$\mathcal{L}(\theta) = \sum_{t=1}^{T} \log(f(y_t | \theta, Y_{t-1})). \tag{10}$$

Then, the covariance matrix of the ML estimator is estimated as the inverse of the negative of the Hessian evaluated at $\widehat{\theta}$.

Thus, assuming that $P_t^{\mathbf{i}}(\widehat{\theta})$ is a smooth transformation of $\widehat{\theta}$ for all t, and that the partial derivatives $\frac{\partial P_t^{\mathbf{i}}(\widehat{\theta})}{\theta_j}$ exist for all $j = 1, 2, \ldots, s$, the delta method applies. Then, the variance of $P_t^{\mathbf{i}}(\widehat{\theta})$ becomes

$$\sigma^2_{P_t^{\mathbf{i}}(\widehat{\theta})} = Var(P_t^{\mathbf{i}}(\widehat{\theta})) = \nabla(P_t^{\mathbf{i}}(\widehat{\theta}))' \left[-\frac{\partial \mathcal{L}(\theta)}{\partial \theta \partial \theta'} |_{\theta = \widehat{\theta}} \right]^{-1} \nabla(P_t^{\mathbf{i}}(\widehat{\theta})), \tag{11}$$

where $\nabla(P_t^{\mathbf{i}}(\widehat{\theta}))$ denotes the gradient of $P_t^{\mathbf{i}}(\theta)$ evaluated at $\widehat{\theta}$. This gradient is obtained by computing $\frac{\partial P_t^{\mathbf{i}}(\theta)}{\partial \theta_s}$, for all $s = 1, 2, \ldots, r$. To this end, if we denote

$$A = \psi_t^{\mathbf{i}}, \tag{12}$$

$$B = \sum_{\mathbf{j}} P_{t-1}^{\mathbf{j}}(\theta) p_{\mathbf{ji}}^*, \tag{13}$$

$$C = \sum_{\mathbf{i}} \psi_t^{\mathbf{i}} \sum_{\mathbf{j}} p_{\mathbf{ji}}^* P_{t-1}^{\mathbf{j}}(\theta), \tag{14}$$

then the probabilities are given by the ratio

$$P_t^{\mathbf{i}}(\theta) = \frac{A \cdot B}{C}. \tag{15}$$

[4] Note that the filter requires an initial condition of $P_0^{\mathbf{i}}(\theta)$, for all $\mathbf{i}$, which typically is the ergodic probability.

Then, the partial derivatives of this ratio are obtained as

$$\frac{\partial P_t^{\mathbf{i}}(\theta)}{\partial \theta_s} = \frac{\left[\frac{\partial A}{\partial \theta_s} B + A \frac{\partial B}{\partial \theta_s}\right] C - AB\frac{\partial C}{\partial \theta_s}}{C^2}, \tag{16}$$

where

$$\frac{\partial A}{\partial \theta_s} = \frac{\partial \psi_t^{\mathbf{i}}}{\partial \theta_s}, \tag{17}$$

$$\frac{\partial B}{\partial \theta_s} = \sum_{\mathbf{j}} \left[\frac{\partial P_{t-1}^{\mathbf{j}}(\theta)}{\partial \theta_s} p_{\mathbf{ji}}^* + P_{t-1}^{\mathbf{j}}(\theta) \frac{\partial p_{\mathbf{ji}}^*}{\partial \theta_s}\right], \tag{18}$$

$$\frac{\partial C}{\partial \theta_s} = \sum_{\mathbf{i}} \left[\frac{\partial \psi_t^{\mathbf{i}}}{\partial \theta_s} \sum_{\mathbf{j}} p_{\mathbf{ji}}^* P_{t-1}^{\mathbf{j}}(\theta) + \psi_t^{\mathbf{i}} \left(\frac{\partial p_{\mathbf{ji}}^*}{\partial \theta_s} P_{t-1}^{\mathbf{j}}(\theta) + p_{\mathbf{ji}}^* \frac{\partial P_{t-1}^{\mathbf{j}}(\theta)}{\partial \theta_s}\right)\right]. \tag{19}$$

Next, we explain the obtained estimation procedure from a practical point of view. The above equations suggest an iterative algorithm for finding the estimate of the gradients $\nabla(P_t^{\mathbf{i}}(\theta))$, $t = 1, \ldots, T$. Starting from the initial condition at $t = 0$ for the filtered probabilities used in the Hamilton's filter and the gradient for each $\mathbf{i}$, $\{P_0^{\mathbf{i}}(\theta), \frac{\partial P_0^{\mathbf{i}}(\theta)}{\partial \theta_1}, \frac{\partial P_0^{\mathbf{i}}(\theta)}{\partial \theta_2}, \ldots, \frac{\partial P_0^{\mathbf{i}}(\theta)}{\partial \theta_r}\}$, which are the magnitudes on the right-hand sides of (16)-(19), the left-hand sides of these expressions then produce a new estimate for the gradient $\nabla(P_1^{\mathbf{i}}(\theta))$.[5] This estimate can be used to re-evaluate (16)-(19), and we can continue iterating in this fashion until the last gradient $\nabla(P_T^{\mathbf{i}}(\theta))$. The gradients can now be used in (11) to obtain the variances of the filtered probabilities.

Finally, the inference about the unobserved state s_t at t, given observations up to t, can be expressed in terms of the probability distribution of the filtered probabilities:

$$P(s_t = i|\theta, Y_t) := P_t^i(\theta) = \sum_{\mathbf{i}, i_0 = i} P(s_t^* = \mathbf{i}|\theta, Y_t), \tag{20}$$

for all $i = 0, 1, 2, \ldots, K-1$. Then, it follows that

$$\nabla(P(s_t = i|\theta, Y_t)) = \sum_{\mathbf{i}, i_0 = i} \nabla(P(s_t^* = \mathbf{i}|\theta, Y_t)) \tag{21}$$

and, therefore, we can estimate the variance of the filtered probabilities $P(s_t = i|\theta, Y_t)$ as:

$$\sigma^2_{P_t^i(\widehat{\theta})} = Var(P_t^i(\widehat{\theta})) = \sum_{\mathbf{i}, i_0 = i} \nabla(P(s_t^* = \mathbf{i}|\theta, Y_t))' \left[-\frac{\partial \mathcal{L}(\theta)}{\partial \theta \partial \theta'}|_{\theta = \widehat{\theta}}\right]^{-1} \sum_{\mathbf{i}, i_0 = i} \nabla(P(s_t^* = \mathbf{i}|\theta, Y_t)). \tag{22}$$

2.3 Variance of smoothed probabilities

When using a time series to separate the K possible states, probability statements about s_t that incorporate the overall information Y_t are frequent in practice. Such probability statements are given by the

[5] In the simulations and in the empirical applications, we use the ergodic probabilities as a initial condition of $P_0^{\mathbf{i}}(\theta)$. For the means, the variances and the autoregressive parameters, the initial conditions for $\frac{\partial P_0^{\mathbf{i}}(\theta)}{\partial \theta_{\mathbf{i}}}$ are zero. In the case of the transition probabilities, $\frac{\partial P_0^{\mathbf{i}}(\theta)}{\partial p_{\mathbf{ij}}}$, this derivative is positive when $\mathbf{i} = \mathbf{j}$, and negative when $\mathbf{i} \neq \mathbf{j}$. This implies that the filtered probability of regime $\mathbf{i}$ increases with $p_{\mathbf{ii}}$, and decreases with $p_{\mathbf{ij}}$, for $\mathbf{i} \neq \mathbf{j}$.

full-sample smoothed probabilities $P(s_t = i|\theta, Y_T) := Sm_t^i(\theta)$. The smoother suggested in Kim (1994) expresses these probabilities in a way that they may be obtained recursively from the filter probabilities:

$$Sm_t^i(\theta) = P_t^i(\theta) \sum_{j=0}^{K-1} \frac{Sm_{t+1}^j(\theta) p_{ij}}{\sum\limits_{r=0}^{K-1} P_t^r(\theta) p_{rj}}, \tag{23}$$

for all $i = 0, 1, 2, \ldots, K-1$. The smoother operates as a backward algorithm, starting from $Sm_T^j(\theta) = P_T^j(\theta)$ and running backwards in time.

Following the same reasoning as in the case of the filtered probabilities, the variance of the smoothed probabilities can be computed as follows:

$$Var(Sm_t^i(\widehat{\theta})) := \sigma^2_{Sm_t^i}(\widehat{\theta}) = \nabla(Sm_t^i(\widehat{\theta}))' \left[-\frac{\partial \mathcal{L}(\theta)}{\partial\theta\partial\theta'}|_{\theta=\widehat{\theta}} \right]^{-1} \nabla(Sm_t^i(\widehat{\theta})). \tag{24}$$

The gradient $\nabla(Sm_t^i(\theta))$ can be obtained by computing $\frac{\partial Sm_t^i(\theta)}{\partial\theta_s}$ as

$$\begin{aligned} \frac{\partial Sm_t^i(\theta)}{\partial\theta_s} &= \frac{\partial P_t^i(\theta)}{\partial\theta_s} \sum_{j=0}^{k-1} \frac{Sm_{t+1}^j(\theta) p_{ij}}{\sum\limits_{r=0}^{k-1} P_t^r(\theta) p_{rj}} + P_t^i(\theta) \times \\ & \sum_{j=0}^{k-1} \frac{\left[\frac{\partial Sm_{t+1}^j(\theta)}{\partial\theta_s} p_{ij} + Sm_{t+1}^j(\theta) \frac{\partial p_{ij}}{\partial\theta_s} \right] \sum\limits_{r=0}^{k-1} P_t^r(\theta) p_{rj} - Sm_{t+1}^j(\theta) p_{iji} \sum\limits_{r=0}^{k-1} \left[\frac{\partial P_t^r(\theta)}{\partial\theta_s} p_{rj} + P_t^r(\theta) \frac{\partial p_{rj}}{\partial\theta_s} \right]}{\left(\sum\limits_{r=0}^{k-1} P_t^r(\theta) p_{rj} \right)^2}, \end{aligned} \tag{25}$$

for $s = 1, 2, ..., r$. Taking into account that $Sm_T^i(\theta) = P_T^i(\theta)$, equation (25) suggests using iterative backwards algorithm to estimate the gradients $\nabla(Sm_t^i(\theta))$, for all $t = T-1, T-2, \ldots, 1$. More specifically, starting at $t = T-1$, $\frac{\partial P_T^z(\theta)}{\partial\theta_s}$, $P_T^z(\theta)$,$P_{T-1}^z(\theta)$, and the transition probabilities are the magnitudes on the right-hand side of (25), producing an estimate of $\frac{\partial Sm_{T-1}^i(\theta)}{\partial\theta_s}$. This estimate, together with $\frac{\partial P_{T-2}^z(\theta)}{\partial\theta_s}$, $P_{T-2}^z(\theta)$,$Sm_{T-1}^z(\theta)$ can now be used to calculate $\frac{\partial Sm_{T-2}^i(\theta)}{\partial\theta_s}$. Then, continue iterating backwards in this way until the first gradient $\nabla(Sm_1^i(\theta))$ is calculated.

2.4 Asymptotic distribution

Dealing with the issue of inference on the filtered and smoothed probabilities requires a distribution for the unobserved probabilities. Using the results developed above, for a given regime i, the filtered probabilities follow asymptotically normal distributions, with means approximated by $P_t^i(\widehat{\theta})$ and $Sm_t^i(\widehat{\theta})$, and variances approximated by $\sigma^2_{P_t^i(\widehat{\theta})}$ and $\sigma^2_{Sm_t^i(\widehat{\theta})}$, respectively.

Therefore, $100(1-\alpha)\%$ confidence intervals have the form

$$CI_\alpha(PR_t^i) = [b_{1-\alpha/2}^{PR_t^i}, b_{\alpha/2}^{PR_t^i}], \tag{26}$$

where PR_t^i is $P_t^i(\widehat{\theta})$ for filtered probabilities and $Sm_t^i(\widehat{\theta})$ for smoothed probabilities, and b_α is defined as $P(N(PR_t^i, \sigma^2_{PR_t^i}) > b_\alpha^{PR_t^i}) = \alpha$. These intervals can be improved on, in the sense that at least 95% of the samples, the estimated value of PR_t^i will be in $\left[\max\left(b_{1-\alpha/2}^{PR_t^i}, 0\right), \min\left(b_{\alpha/2}^{PR_t^i}, 1\right)\right]$.[6]

[6] An alternative would be to approximate the distribution of the filtered and smoothed probabilities with a Beta distribution with means $P_t^i(\widehat{\theta})$ and $Sm_t^i(\widehat{\theta})$, and variances $\sigma^2_{P_t^i(\widehat{\theta})}$ and $\sigma^2_{Sm_t^i(\widehat{\theta})}$, respectively.

3 Monte Carlo simulation

3.1 Inference on turning point dates

To evaluate the impact of the model parameters on the uncertainty associated to filtered and smoothed probabilities from (1), we follow Cavicchioli (2014) to consider the univariate Markov-switching model

$$y_t = \mu_{s_t} + \phi_{s_t}(y_{t-1} - \mu_{s_{t-1}}) + \varepsilon_t, \tag{27}$$

where $\varepsilon_t \sim N(0, \sigma^2_{s_t})$.[7] In this model, s_t is a two-state Markov-chain process of order one that takes the value 0 in the first regime, and 1 in the second regime.

Using this model, we generate dummy variables s_t of zeroes and ones of length $T = 250$, which are used to simulate different sequences of the two regimes of the unobserved state variable, assuming that $p_{00} = p_{11} = 0.9$. In addition, we generate shocks ε_t with variances $\sigma_0^2 = 0.5$ and $\sigma_1^2 = 1$. The dynamics of the generated time series y_t are assumed to follow an autoregressive process of order one, with autoregressive parameters $\phi_0 = 0.2$ and $\phi_1 = 0.5$, and within-state means $\mu_0 = 1$ and $\mu_1 = -1$, respectively.

The estimated probabilities of state 1, the standard deviation of the point estimates, and the dates for which the state variable takes the value one in this simulation (shaded areas) are plotted in Figure 1. The first and the third panels of this figure, which plot the filtered and smoothed probabilities respectively, show that the model accurately captures the two different regimes, because the probability of state 1 is high when $s_t = 1$, and low when $s_t = 0$. The second and the fourth panels in the figure show that the standard deviation of the probability of state 1 in the middle of the two phases is low, which implies that the information content of the filtered probabilities as classification rules is high in this example.

Notably, the uncertainty surrounding the point estimates of the probabilities increases significantly around the dates on which the model detects regime changes. As a result, the confidence intervals become wide around the turning points, indicating a high degree of uncertainty in identifying the state at these dates. This suggests that rapid jumps in the uncertainty associated with a particular regime might be useful in determining the timing of the turning points from (or towards) this regime.

3.2 Sensitivity analysis

In this section, we set up several Monte Carlo experiments to examine how the model parameters might affect the measures of the uncertainty of the filtered and smoothed probabilities. For this purpose, Figures 2 and 3 show the medians over the T generated observations of the standard deviation of the probabilities of state 1 when the baseline parameters change as follows: in Panel A, $\mu_0 - \mu_1 = 1, 1.1, ..., 3$; in Panel B, $\sigma_0^2 = \sigma_1^2 = 0.5, 0.7, ..., 2$; in Panel C, $p_{00} = p_{11} = 0.5, 0.6, ..., 0.9$; and in Panel D, $\phi_0 = \phi_1 =$

[7] Note that the model implies a general class of Markov-switching dynamics: it allows the regimes affecting the mean, the autoregressive parameters, and variance. Without loss of generalization, the lag length is restricted to one.

$0.1, 0.2, ..., 0, 9$. The figures show that the uncertainty over the inference on the computed filtered and smoothed probabilities increases with the variance of the shocks, which means that it is more difficult to distinguish the regime switches in processes that include large shocks.

In addition, the uncertainty of the filtered and smoothed probabilities decreases with the difference of within-state means because the regimes are clearly separated from one another, which implies that the changes in regime are clear from the data. The uncertainty also decreases with the inertia of the states, and with the autoregressive parameter of the process because, in both cases, the regimes are persistent and the number of turning points, which are the source of the uncertainty, diminishes.

4 Empirical illustrations

The empirical relevance of the proposed theory is illustrated through two applications. The first application focuses on a formal statistical model of business cycle phase shifts, which has probably been the most extensive application of Markov-switching autoregressive models. In the second application, three states are allowed to characterize the infrequent changes in the mean and variance of the ex-post real interest rate.

4.1 Analysis of business cycle regimes

Hamilton (1989) proposes that output growth may follow one of two different autoregressions, with high and low means, depending on whether the economy is expanding or contracting, with the shift between the regimes governed by the outcome of an unobserved first-order Markov chain.[8] Accordingly, the growth rate of U.S. quarterly real GDP from 1951.1 to 2016.3, y_t, is allowed to switch according to

$$y_t = \mu_{s_t} + \varepsilon_t. \tag{28}$$

At time t, we label $s_t = 0$ as expansions and $s_t = 1$ as recessions, with $\mu_1 < \mu_0$.[9] Deviations from this mean growth rate are created by ε_t, which is an *i.i.d.* Gaussian stochastic disturbance, with a mean of zero and variance σ^2. Therefore, GDP is expected to exhibit high (usually positive) growth rates in expansions and low (usually negative) growth rates in recessions. This pattern is depicted in the first panel of Figure 4, which plots the growth rate of U.S. GDP, along with shaded areas that represent the NBER recessions.

The maximum likelihood estimates, reported in Table 1, show that the transition probabilities are highly persistent ($\widehat{p}_{00} = 0.95$, $\widehat{p}_{11} = 0.69$), and that the within-state means are separate from each other ($\widehat{\mu}_0 = 0.96$, $\widehat{\mu}_1 = -0.48$). According to our simulation results, the separated within-state means and the

[8] Camacho and Perez-Quiros (2007) show that a model that decomposes output growth into a state-dependent mean and a stationary process captures the U.S. business cycle dynamics with high precision.

[9] Based on Chauvet and Hamilton (2006), this is the model used by the Federal Reserve Bank of St. Louis to compute the GDP-based recession indicator.

persistence of the states help the Markov-switching autoregressive model to compute accurate inferences of the U.S. business cycle dates. The ability of this simple model to characterize the U.S. business cycle features is shown in the last two panels of Figure 4, which plot the filtered and smoothed probabilities of a low-growth regime. The figure show that the official business cycle dates correspond fairly closely to the inferences about the unobservable state variable.[10]

In some cases, analysts might be interested in converting filtered and smoothed probabilities into specific sets of dates, establishing the timing of shifts between business cycle phases or turning points. This requires specific rules to establish whether a particular quarter was an expansion quarter or a recession quarter. Hamilton (1989) suggests that a natural metric might be based on whether analysts can conclude that the economy is more likely to be in a recession, in other words, $P_t^i(\widehat{\theta}) > 0.5$ or $Sm_t^i(\widehat{\theta}) > 0.5$.[11] If we denote a business cycle turning point as having occurred the filtered probabilities move from below or above 0.5, the dates identified by this metric roughly coincide with the NBER in most cases.

In spite of the simplicity of this method, using the filtered probabilities results in a few false negatives in 1970.3 and 1974.2, and in the 2001 recession. Using the smoothed probabilities, the false negatives refer to the 1970 and 2001 recessions.[12] Now, we need to establish whether the uncertainty surrounding the point estimates may help to identify the actual business cycle phase on these dates.

In line with our Monte Carlo results, Figure 5 shows that the standard deviations of the filtered and smoothed probabilities are low, with average figures of about 0.04 in both cases. However, the standard deviations exhibit sudden jumps near the business cycle turning points. Of particular interest are the sudden jumps in variance exhibited in the 1970 and the 2001 recessions. Using the filtered probabilities, the standard deviation rises from 0.02 in 1969.3 to 0.14 in 1969.4 (the NBER peak), and falls from 0.18 in 1970.4 (NBER trough) to almost 0 in 1971.1. In addition, it rises from 0.03 in 2000.4 to 0.14 in 2001.1 (the NBER peak), and falls from 0.20 in 2001.4 (NBER trough) to 0.06 in 2002.1. Using the smoothed probabilities, the standard deviation rises from 0.14 in 1969.3 to 0.28 in 1969.4 (the NBER peak), and falls from 0.24 in 1970.4 (NBER trough) to almost 0 in 1971.1. In addition, it rises from 0.10 in 2000.4 to 0.20 in 2001.1 (the NBER peak), and falls from 0.12 in 2001.4 (NBER trough) to 0.04 in 2002.1. This illustrates the usefulness of using the uncertainty surrounding the point estimates of the filtered and smoothed probabilities to establish the set of business cycle turning points.

An alternative way of dealing with the uncertainty of the filtered and smoothed probabilities of a recession is to use the confidence intervals, as shown in Figure 6. The figure shows that when a recession begins, the amplitude of the estimated confidence intervals increases, showing an increase in the uncertainty of identifying the business cycle phase. When a recession consolidates, the amplitude of the

[10] Note that for model (28), the formulas for the filtered probabilities and the derivatives forming the gradients can be obtained by assuming that $\psi_t^{\mathbf{i}}$ is the pdf of a Gaussian distribution with mean μ_{i_0} and variance σ^2, where $i_0 \in \{0, 1\}$.

[11] Chauvet and Hamilton (2006) and Chauvet and Piger (2008) suggest other metrics, also based on a comparison between filtered and/or smoothed probabilities with pre-specified thresholds.

[12] The difficulties of identifying this short and mild recessions are documented in Kliesen (2003) and Hamilton (2011).

estimated confidence intervals reduces, and the lower limits of the confidence intervals take values close to 1. As the recession period ends, the amplitude of the confidence interval increases again, which could be interpreted as a signal of a change in state.

One final remark is the usefulness of having the distribution of the single inferences of both the filtered and smoothed probabilities in real time. Computing inferences in real time differs from the framework shown so far because the GDP figures, as originally released by the Bureau of Economic Analysis, can differ substantially from the historical series now available. To overcome this potential drawback, we use the real-time data set archived at the Federal Reserve Bank of Philadelphia, which includes the history of GDP values that would actually have been available to a researcher at any given point in time.[13]

Using these time series, the inferences are computed in a recursive way. The first exercise begins in 1956.4, with data available from 1951.1 to 1956.3, from which we estimate the Markov-switching autoregressive model and collect the filtered probability of a recession for 1956.3, and its variance. Then, the sample is updated with the time series available in 1966.1 (data from 1951.1 to 1956.4), for which the model is re-estimated, and we collect the inference for 1956.4 and its variance. The procedure continues iteratively until the final time series available in 2016.4, which produces the filtered probability of a recession in 2016.3 and its corresponding variance.

Figure 7 displays the filtered probabilities of recessions and their confidence intervals, as they were inferred in real time. As expected, the NBER business cycle becomes more difficult to identify in real time. Again, following the 0.5 rule, some false positives appear, for example in 1967.1, 1979.2, and the year after the 1990 recession. For their potential damage, the time delays with which the last two recessions are identified are of special interest in this analysis. The probability of a recession does not cross the 0.5 line until 2001.3 and 2008.4. However, the threshold belongs to the confidence intervals in 2001.2 and 2008.1, the first quarters of their respective NBER recessions.

One final remark of the real-time analysis concerns the increased volatility of the inferred probability of a recession that appears since the mid-1980s, which corresponds to the Great Moderation period.[14] Kim and Nelson (1999) showed that the volatility reduction in output growth implied a narrowing gap between growth rates during recessions and expansions. In line with our simulation results, this would imply that the probabilities of recessions are estimated with higher uncertainty. Accordingly, performing business cycle inferences became a more difficult challenge, because the business cycle phases are more difficult to identify accurately.

4.2 Analysis of real interest rates regimes

Garcia and Perron (1996) employ a Markov-switching model to account for regime switches in an autoregressive model of U.S. ex-post real interest rate. In particular, they propose the following process with a

[13] We fill in some minor gaps in this series, following Chauvet and Hamilton (2006).

[14] In independent contributions, Kim and Nelson (1999) and McConnell and Perez-Quiros (2000) detected a substantial moderation in output growth volatility, showing that this moderation is well modeled as a single break in the mid-1980s.

three-state Markov-switching mean and variance:

$$r_t = \mu_{s_t} + \phi_1 \left(r_{t-1} - \mu_{s_{t-1}} \right) + \phi_2 \left(r_{t-2} - \mu_{s_{t-2}} \right) + \varepsilon_t, \tag{29}$$

where ε_t is an *i.i.d.* Gaussian stochastic disturbance with a mean of zero and variance $\sigma^2_{s_t}$. The authors propose modeling s_t as the outcome of an unobserved three-state, first-order Markov process. In this context, s_t can take the value 0 in the state they label as a low mean, 1 in the state they label as a middle mean, or 2 in the state they label as a high mean.

Our empirical example uses the quarterly ex-post real interest rate, which is the difference between the end-of-quarter figures of the nominal interest rate and the inflation rate, using data from 1934.1 to 2016.3.[15] The time series, shown in Figure 8, exhibits occasional jumps across negative, positive, and highly positive values. The real rates became negative (and highly volatile) until the 1950s, during the inflationary period of the 1970s, and since the beginning of the new century. They reached (low) positive values in the 1950s, 1960s, early 1970s, and in the late 1980s and 1990s. Finally, the early 1980s saw an unprecedented huge spike in interest rates.

The ML estimation results are reported in Table 2. The within-regime estimated means suggest that the dynamics of the time series are characterized by three distinct phases of low, middle, and high levels of real interest rates. The volatility of the interest rate is significantly higher in the low regime, while it is closer and smaller in the middle and high regimes. Notably, the autoregressive parameters are close to zero. In line with Garcia and Perron (1996), this suggests that the real interest rate may not have a unit root because most of its persistence is accounted for by the regime switching in the mean. In addition, the three regimes are highly persistent because the estimated probability that the same regime prevails is above 0.95 in the three regimes. Finally, the smooth behavior of the estimated ex-ante real interest rate, shown in Figure 8, helps to see the interest rate as a constant subject that shifts across the three regimes.

Figure 9 shows the filtered probabilities when the interest rate is in the low regime (top panel), the middle regime (middle panel), and the high regime (bottom panel), along with their respective 95% confidence intervals.[16] According to our simulation results, the clearly separated within-regime means and the high persistence of the regimes diminish the uncertainty of the estimated probabilities, which leads to very narrow confidence intervals. The figure shows that the confidence intervals are very useful to establish the degree of uncertainty in converting the filtered probabilities into a specific set of dates that provide the timing of shifts across the different interest rate regimes. Establishing the amount of uncertainty over the timing of the turning points is of extreme importance to discriminating among alternative potential explanations offered for the shifts in the real interest rate.

According to Figure 9, the dating of the first shift from the low regime to the middle regime, and the

[15] The nominal interest rate is the three-month treasury bill rate (TB3MS), and the inflation rate is constructed from the CPI (CPIAUCNS) at annual rates, non-seasonally adjusted.

[16] The full sample smoothed probabilities were also computed, but they differed very little from the filtered probabilities.

subsequent return to the low regime, seem to be associated with single distinct breaks at about 1953.3 and 1973.4, respectively. In contrast, the dates of the regime shifts that have occurred since this date are more uncertain. Although the filtered probability indicates that the shift from the low regime to the high regime occurred in 1981.2, and that the shift to the middle regime occurred in 1986.4, the confidence intervals suggest that the shifts could occur in 1980.4 and 1990.2, respectively.[17] Finally, the cut in the cost of borrowing that characterizes the beginning of the new century is dated as 2003.1 by the filtered probabilities, while the confidence intervals suggest that the shift could have occurred earlier, in 2002.1.

5 Conclusions

We contribute to the rapid growth in the literature on Markov-switching autoregressive models in the last three decades by developing a distribution theory useful for testing and computing inferences from filtered and smoothed probabilities. The proposed framework is flexible enough to cover a large part of the literature on Markov-switching autoregressive models.

Using a Monte Carlo analysis, we showed that the uncertainty over the probabilities diminishes when the states are separated, the variance of the shocks is low, and the time series or the regimes are persistent. Interestingly, the variance of the probabilities increases dramatically about the dates of phase changes. Therefore, we consider that the confidence intervals of the probabilities could be used, along with the point estimates, to identify changes in the regime of the unobserved state variable.

We show the empirical relevance of using the confidence intervals of filtered and smoothed probabilities computed within the Markov-switching autoregressive models framework by capturing the nonlinearities in the U.S. business cycles of quarterly GDP growth rates. In particular, the confidence intervals are useful to identifying the U.S. recessions when the point estimates of the probabilities are close to the threshold marking the phase changes. In addition, we show that the narrowing gap between the growth rates during the recessions and expansions since the Great Moderation have made the business cycle phases more difficult to identify. Finally, we show that computing confidence intervals could also be useful to identifying the dates of jumps in the regimes of the U.S. real interest rate series, which are important in light of the alternative explanations offered for the regime changes.

[17] This result agrees with the analysis of Garcia and Perron (1996) and contrasts with the study of Huizinga and Mishkin (1986).

References

[1] Ang, A. and Bekaert, G. 2002. Regime switches in interest rates. *Journal of Business & Economic Statistics* 20: 163-182.

[2] Bickel, P., Ritov, Y., and Ryden, T. 1998. Asymptotic normality of the maximum-likelihood estimator for general hidden Markov models. *The Annals of Statistics* 26: 1614-1635.

[3] Camacho, M., and Perez-Quiros, G. 2007. Jump-and-rest effect of U.S. business cycles. *Studies in Nonlinear Dynamics and Econometrics* 11(4): article 3.

[4] Cavicchioli, M. 2014. Analysis of the likelihood function for Markov-switching VAR(CH) models. *Journal of Time Series Analysis* 35: 624-639.

[5] Chauvet, M., and Hamilton, J. 2006. Dating business cycle turning points in real time. In *Nonlinear Time Series Analysis of Business Cycles*, eds. C. Milas, P. Rothman, and D. Van Dijk. Amsterdam: Elsevier Science, pp. 1-54.

[6] Chauvet, M., and Piger, J. 2008. A comparison of the real-time performance of business cycle dating methods. *Journal of Business and Economic Statistics* 26: 42-49.

[7] Dahlquist, M., and Gray, S. 2000. Regime-switching and interest rates in the European monetary system. *Journal of International Economics* 50: 399-419.

[8] Engel, C. 1994. Can the Markov switching model forecast exchange rates? *Journal of International Economics* 36: 151-165.

[9] Engel, C., and Hamilton, J. 1990. Long swings in the dollar: Are they in the data and do markets know it? *American Economic Review* 80: 689-713.

[10] Francq, C., and Zakoïan, J. 2001. Stationarity of multivariate Markov-switching ARMA models. *Journal of Econometrics* 102: 339-364.

[11] Fruhwirth-Schnatter S. 2006. Finite Mixture and Markov Switching Models. Springer: Berlin.

[12] Garcia, R., and Perron, P. 1996. An analysis of the real interest rate under regime shifts. *Review of Economics and Statistics* 78: 111-125.

[13] Hamilton, J. 1989. A new approach to the economic analysis of nonstationary time series and the business cycles. *Econometrica* 57: 357-384.

[14] Hamilton, J. 2011. Calling recessions in real time. *International Journal of Forecasting* 27: 1006-1026.

[15] Hamilton, J. and Lin, G. 1996. Stock market volatility and the business cycle. *Journal of Applied Econometrics* 11: 573-593.

[16] Huizinga, J., and Mishkin, F. 1986. Monetary Policy Regime Shifts and the Unusual Behavior of the Real Interest Rates. *Carnegie-Rochester Conference on Public Policy* 24: 231-274.

[17] Kim, C. 1994. Dynamic linear models with Markov-switching. *Journal of Econometrics* 60: 1-22.

[18] Kim, C., and Nelson, C. 1999. Has the U.S. economy become more stable? A Bayesian approach based on a Markov-switching model of the business cycle. *The Review of Economics and Statistics* 8: 608-616.

[19] Kim, C., Nelson, C. and Startz, R. 1998. Testing for Mean Reversion in Heteroskedastic Data Based on Gibbs-Sampling-Augmented Randomization. *Journal of Empirical Finance* 5: 131-154.

[20] Kliesen, K. 2003. The 2001 recession: How was it different and what developments may have caused it? *Federal Reserve Bank of St. Louis Review* September/October: 23-37.

[21] McConnell, M., and Perez-Quiros, G. 2000. Output fluctuations in the United States: What has changed since the early 1980s? *American Economic Review* 90: 1464-1476.

[22] Pagan, A., and Schwert, G. 1990. Alternative models for conditional stock volatility. *Journal of econometrics* 45: 267-290.

Table 1. MS model for GDP

μ_0	μ_1	σ^2	p_{00}	p_{11}
0.96	-0.48	0.57	0.95	0.69
(0.07)	(0.30)	(0.06)	(0.02)	(0.11)

Notes. The estimated model is $y_t = \mu_{s_t} + \varepsilon_t$, where y_t is the rate of growth of GDP, s_t is a two-state unobservable state variable that governs the business cycle dynamics, $\varepsilon_t \sim iidN(0,\sigma^2)$, and $p_{ii} = p\left(s_t = i/s_{t-1} = i\right)$. Standard errors are in parentheses.

Table 2. MS model for interest rate

		μ_0	μ_1	μ_2	ϕ_1	ϕ_2		
		-2.01	1.50	5.03	0.19	0.04		
		(0.55)	(0.19)	(0.65)	(0.06)	(0.06)		
σ_0^2	σ_1^2	σ_2^2	p_{00}	p_{01}	p_{10}	p_{11}	p_{20}	p_{21}
27.56	2.92	5.34	0.9886	0.0058	0.0174	0.9826	0.0002	0.0440
(3.16)	(0.38)	(1.59)	(0.0085)	(0.0065)	(0.0109)	(0.0109)	(0.0044)	(0.0436)

Notes. The estimated model is $r_t = \mu_{s_t} + \phi_1\left(r_{t-1} - \mu_{s_{t-1}}\right) + \phi_2\left(r_{t-2} - \mu_{s_{t-2}}\right) + \varepsilon_t$, where r_t is the quarterly ex-post real interest rate, s_t is a three-state unobservable state variable, $\varepsilon_t \sim iidN(0,\sigma_{s_t}^2)$, and $p\left(s_t = j/s_{t-1} = i\right) = p_{ij}$. Standard errors are in parentheses.

Figure 1. Simulation

Panel A. Filtered probabilities of state 1

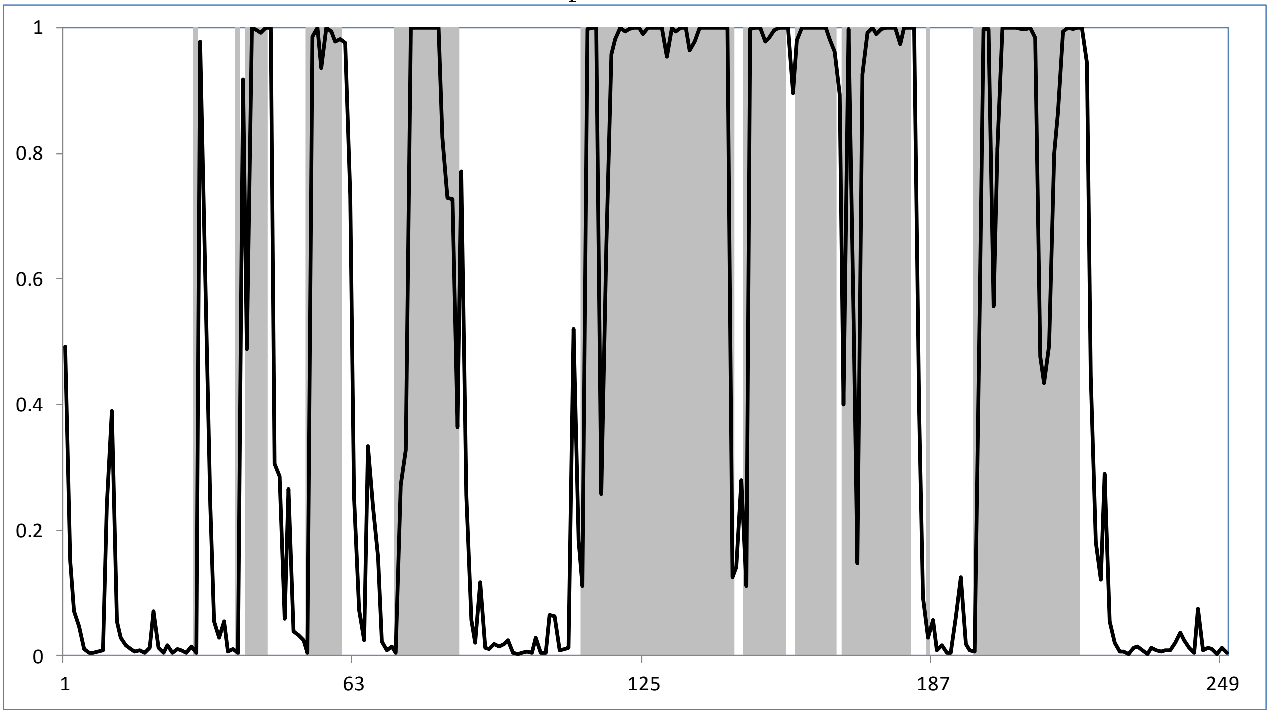

Panel B. Standard deviation of filtered probabilities of state 1

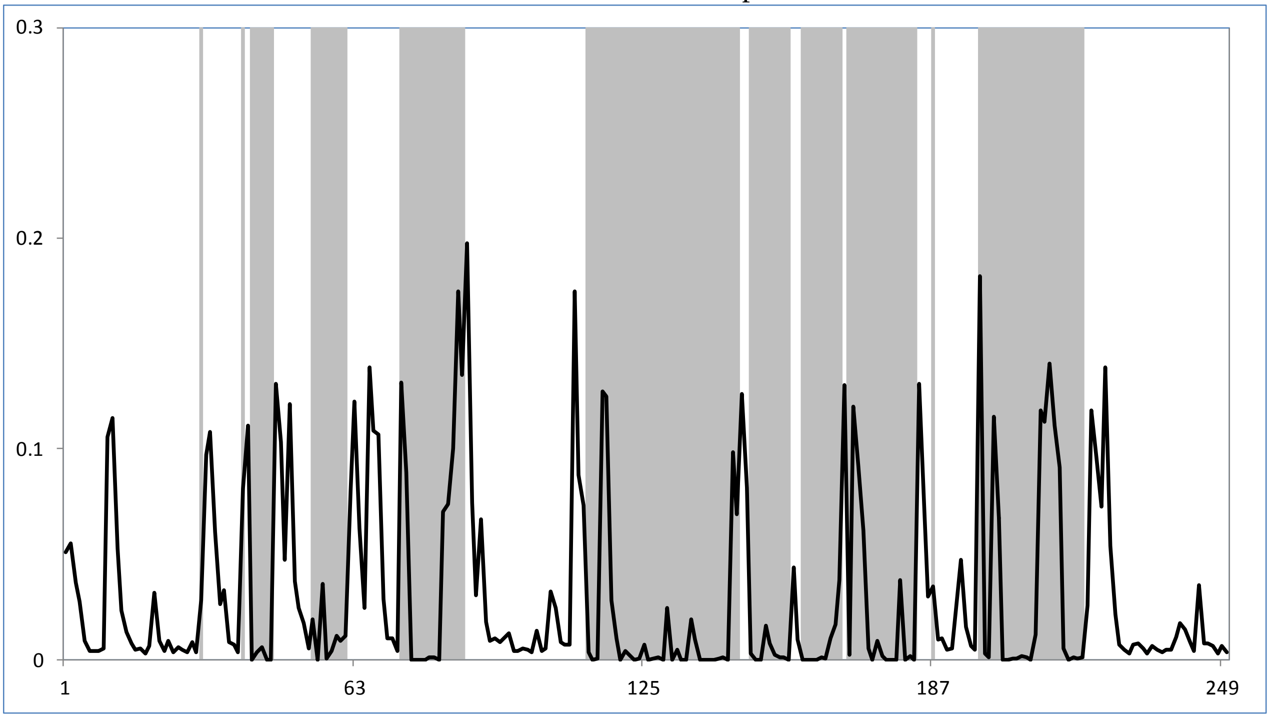

Panel C. Smoothed probabilities of state 1

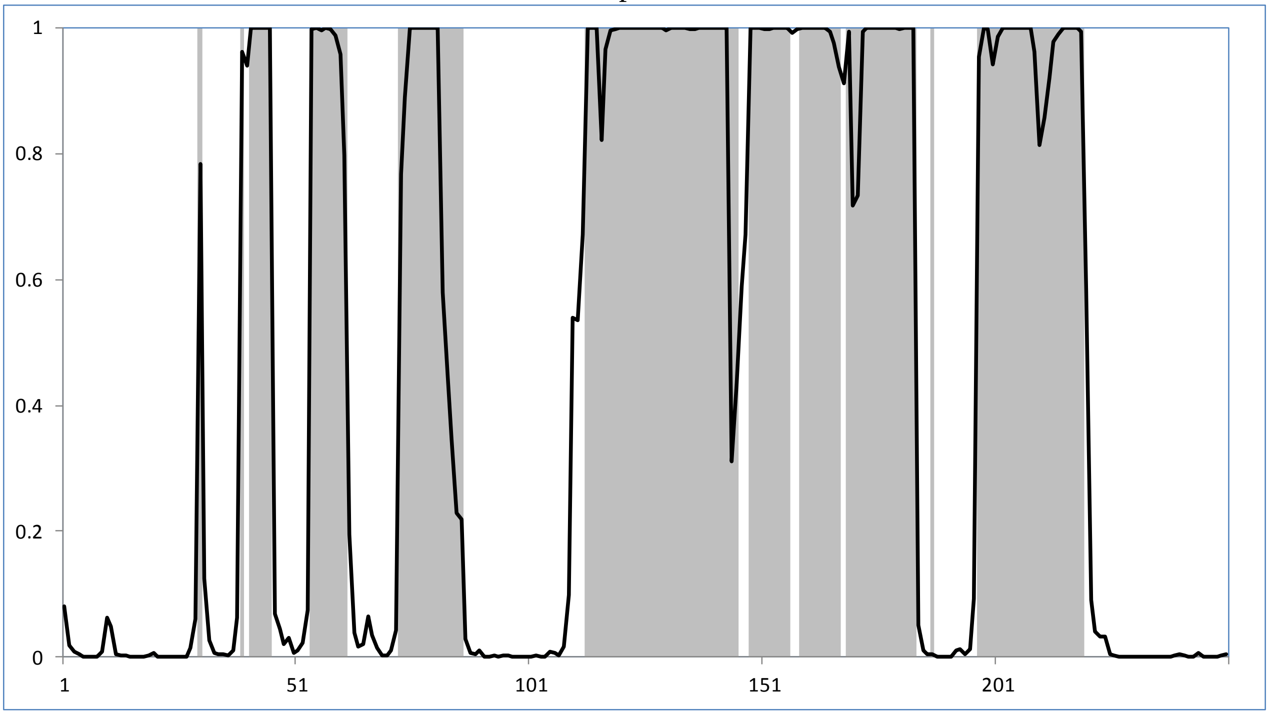

Panel D. Standard deviation of smoothed probabilities of state 1

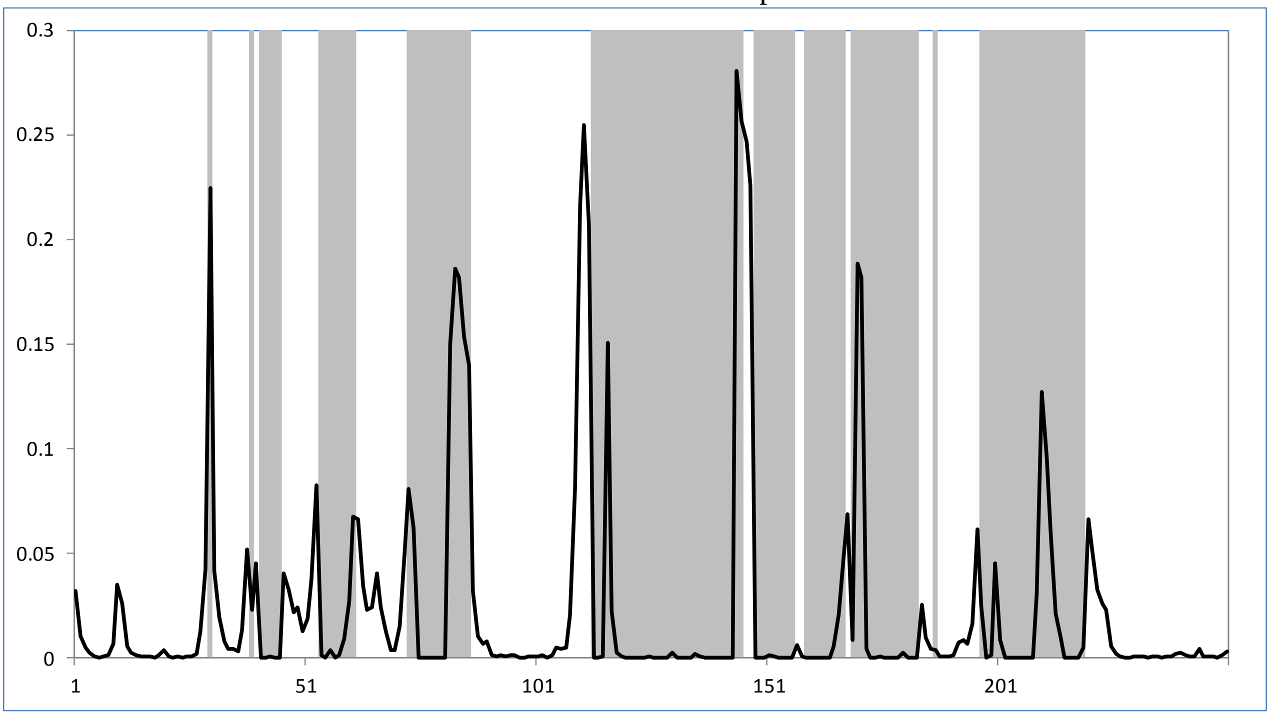

Note. Shaded areas refer to dates where $s_t = 1$. The parameters in the model are $\mu_0 = 1$, $\mu_1 = -1$, $p_{00} = 0.9$, $p_{11} = 0.9$, $\phi_0 = 0.2$, $\phi_1 = 0.5$, $\sigma_0^2 = 0.5$ and $\sigma_1^2 = 1$.

Figure 2. Model parameter changes and standard deviation of filtered probabilities

Panel A. Effect of the within-state means

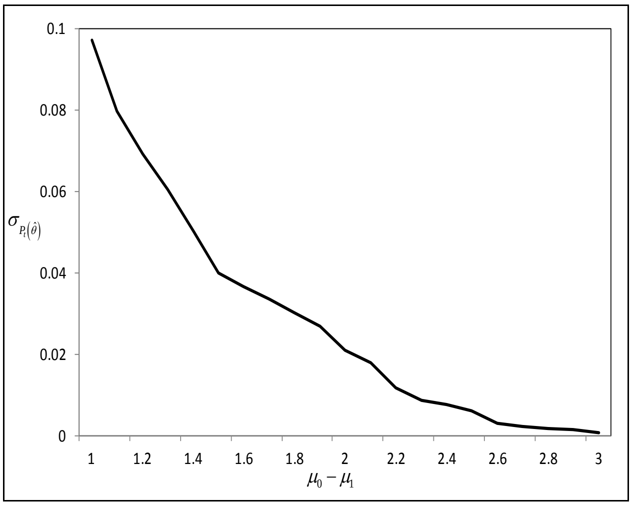

Panel B. Effect of the variance of shocks

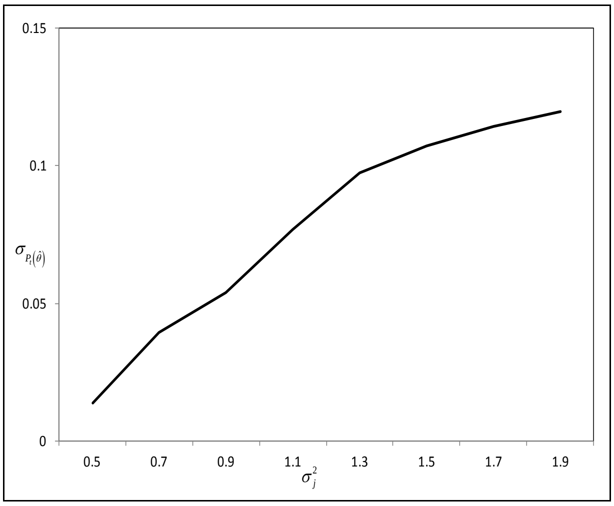

Panel C. Effect of the inertia of the states

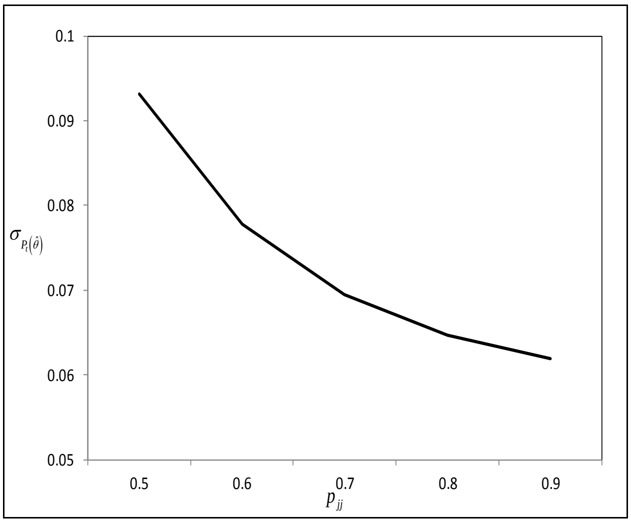

Panel D. Effect of the autoregressive parameter

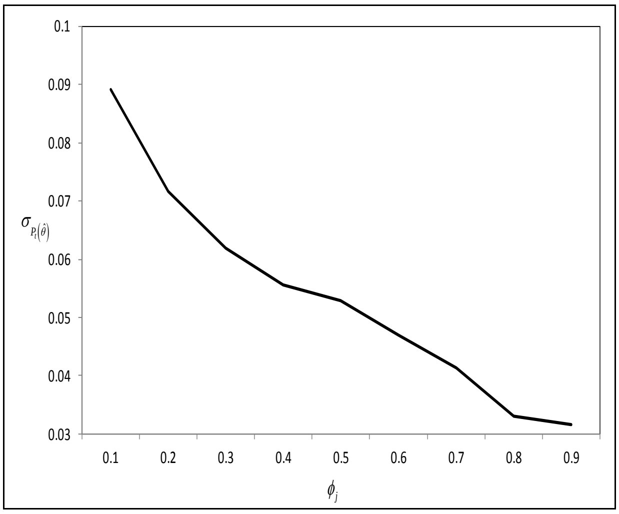

Notes. The parameters in the baseline model are $\mu_0 = 0$, $\mu_1 = -1$, $p_{00} = 0.9$, $p_{11} = 0.9$, $\phi_0 = 0.2$, $\phi_1 = 0.5$, $\sigma_0^2 = 0.5$ and $\sigma_1^2 = 1$. The panels plot the median over the *T* observations of the standard deviations of the filtered probabilities of state 1, $\sigma_{P_t(\hat{\theta})}$, when the baseline parameters change as follows: in Panel A, $\mu_0 - \mu_1 = j$, *j*=1,1.1,1,2,...,3; in Panel B $\sigma_0^2 = \sigma_1^2 = 0,5,0,7,...,2$; in Panel C, $p_{00} = p_{11} = 0.5, 0.6, ..., 0.9$; and in Panel D, $\phi_0 = \phi_1 = 0.1, 0.2, ..., 0.9$.

Figure 3. Model parameter changes and standard deviation of smoothed probabilities

Panel A. Effect of the within-state means

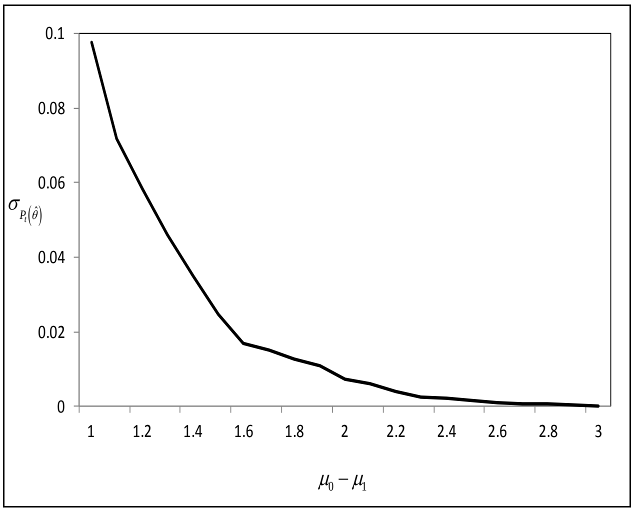

Panel B. Effect of the variance of shocks

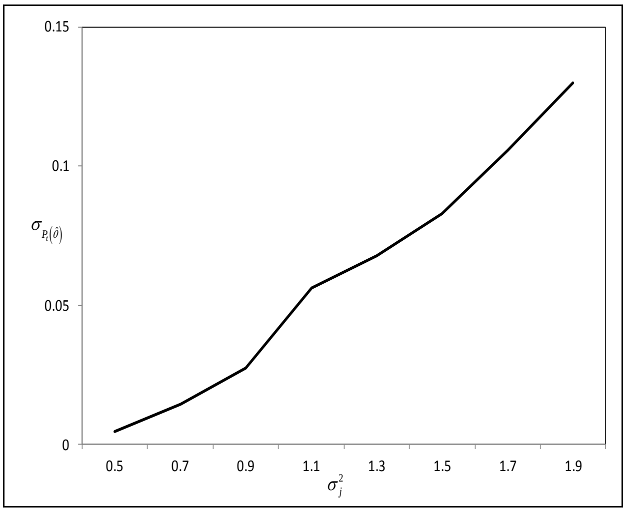

Panel C. Effect of the inertia of the states

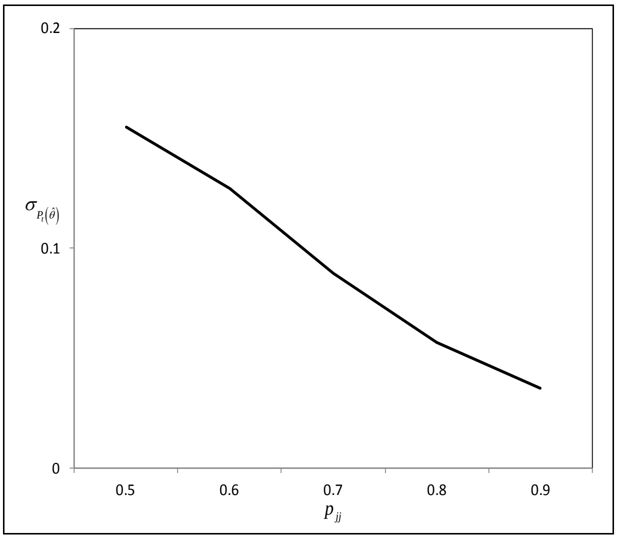

Panel D. Effect of the autoregressive parameter

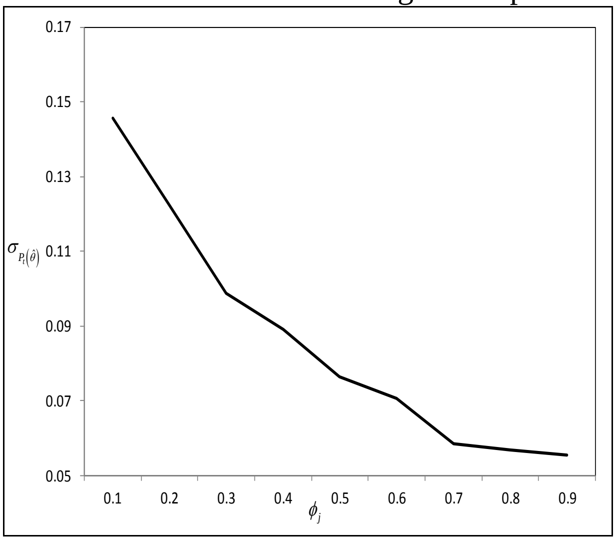

Notes. The parameters in the baseline model are $\mu_0 = 0$, $\mu_1 = -1$, $p_{00} = 0.9$, $p_{11} = 0.9$, $\phi_0 = 0.2$, $\phi_1 = 0.5$, $\sigma_0^2 = 0.5$ and $\sigma_1^2 = 1$. The panels plot the median over the *T* observations of the standard deviations of the smoothed probabilities of state 1 when the baseline parameters change as follows: in Panel A, $\mu_0 - \mu_1 = j$, *j*=1,1.1,1,2,...,3; in Panel B, $\sigma_0^2 = \sigma_1^2 = 0,5,0,7,...,2$; in Panel C, $p_{00} = p_{11} = 0.5, 0.6, ..., 0.9$; and in Panel D, $\phi_0 = \phi_1 = 0.1, 0.2, ..., 0.9$.

Figure 4. In-sample GDP, filtered and smoothed probabilities

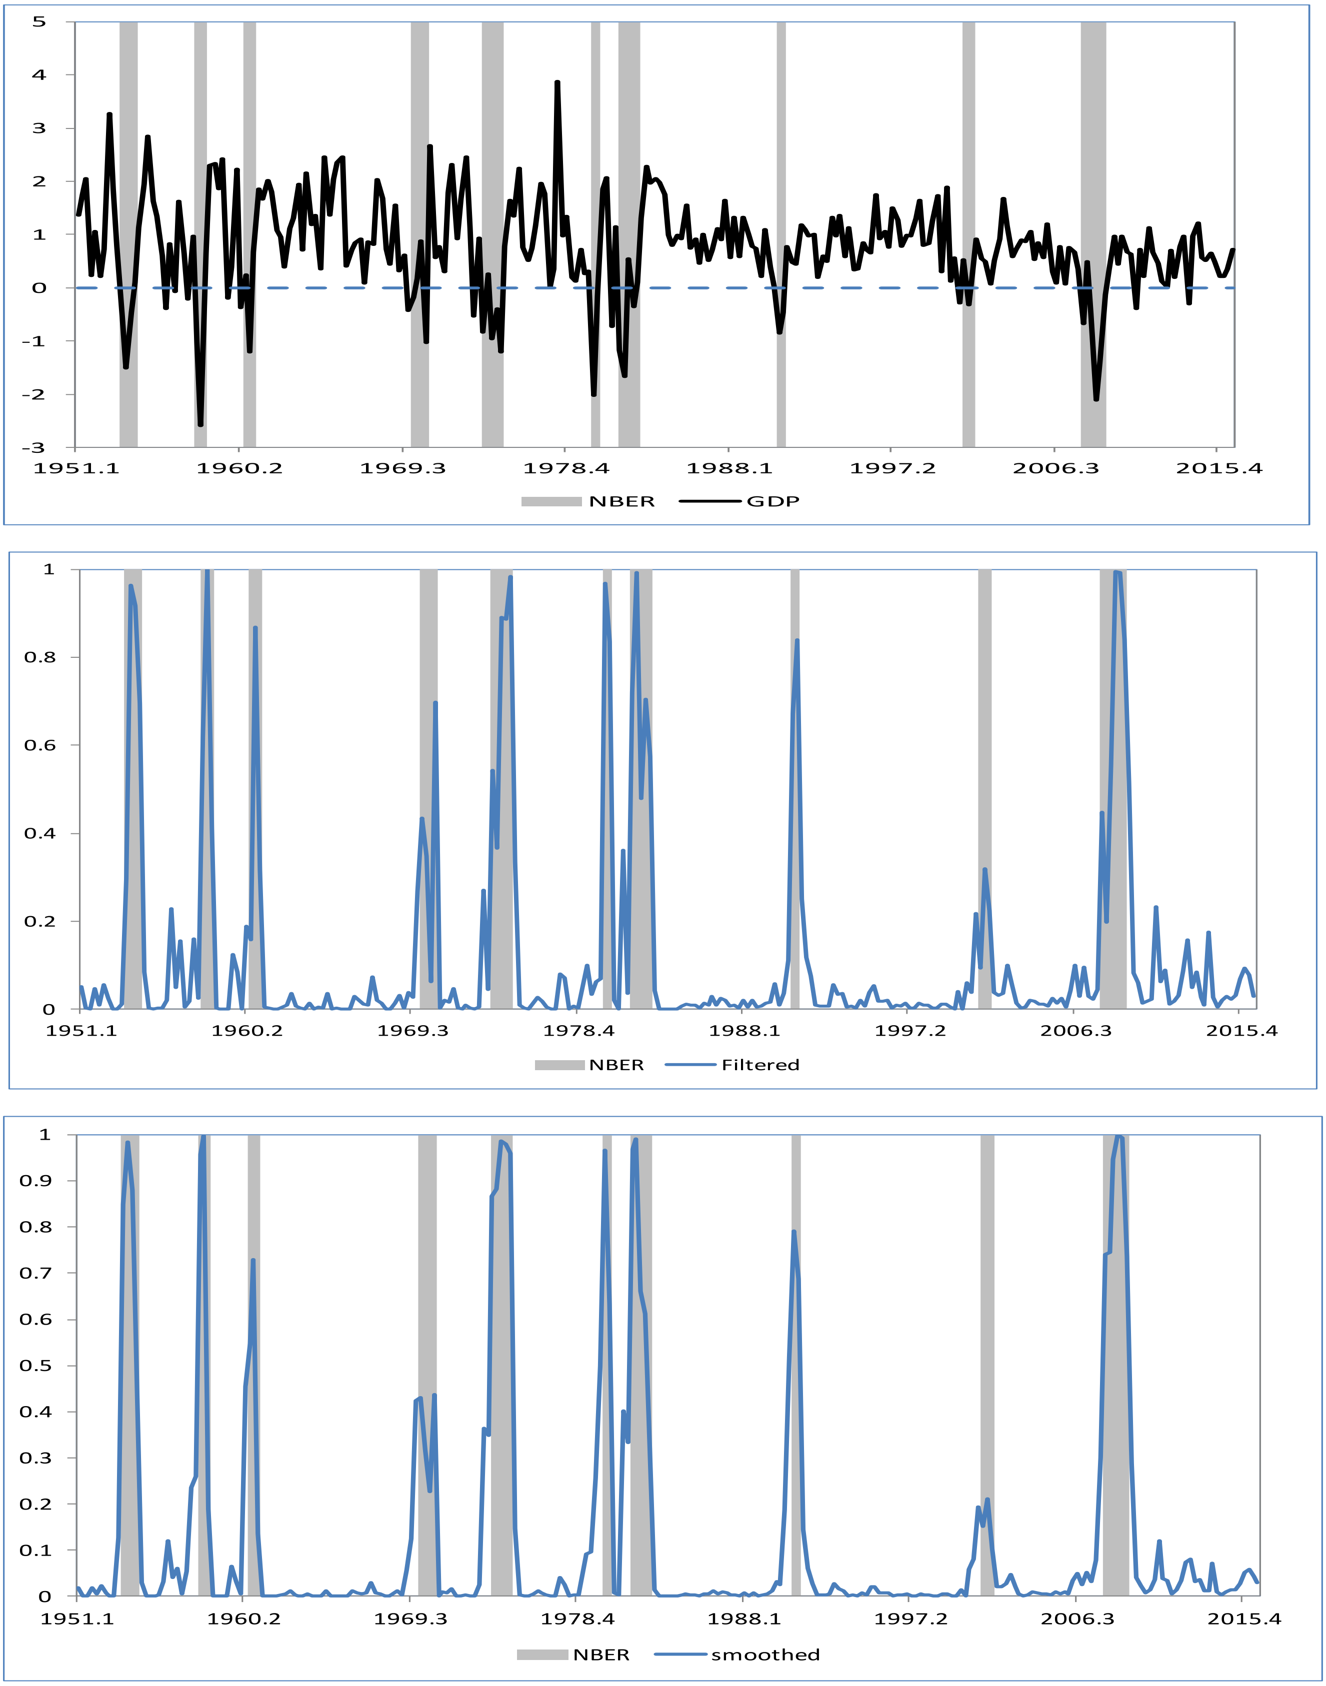

Notes. GDP refers to the growth rate of US quarterly real Gross Domestic Product from 1951.1 to 2016.3 Middle (bottom) chart refers to the inferred filtered (smoothed) probabilities of a regime of low growth. Shaded areas refer to the NBER recessions.

Figure 5. In-sample standard deviations of filtered and smoothed probabilities

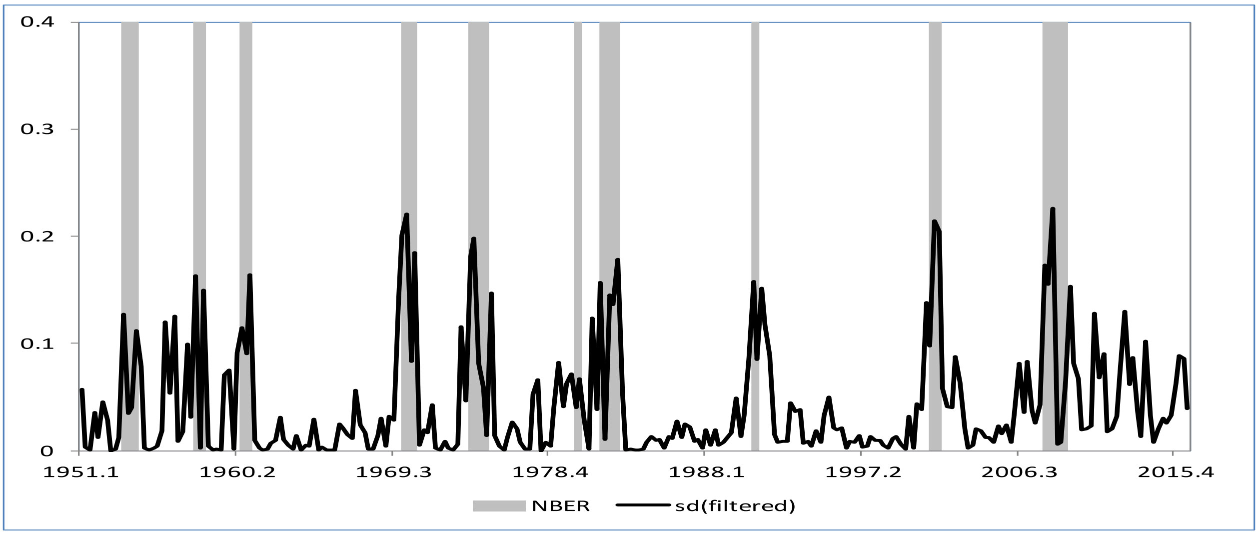

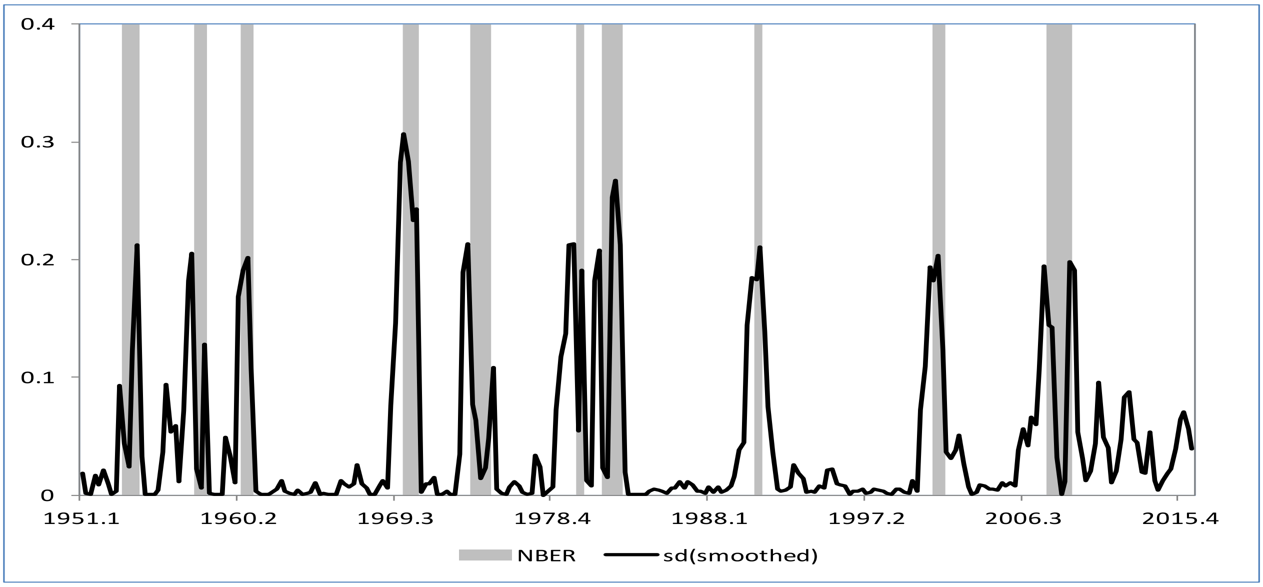

Notes. Shaded areas refer to the NBER recessions.

Figure 6. In-sample confidence intervals

Filtered probabilities

Smoothed probabilities

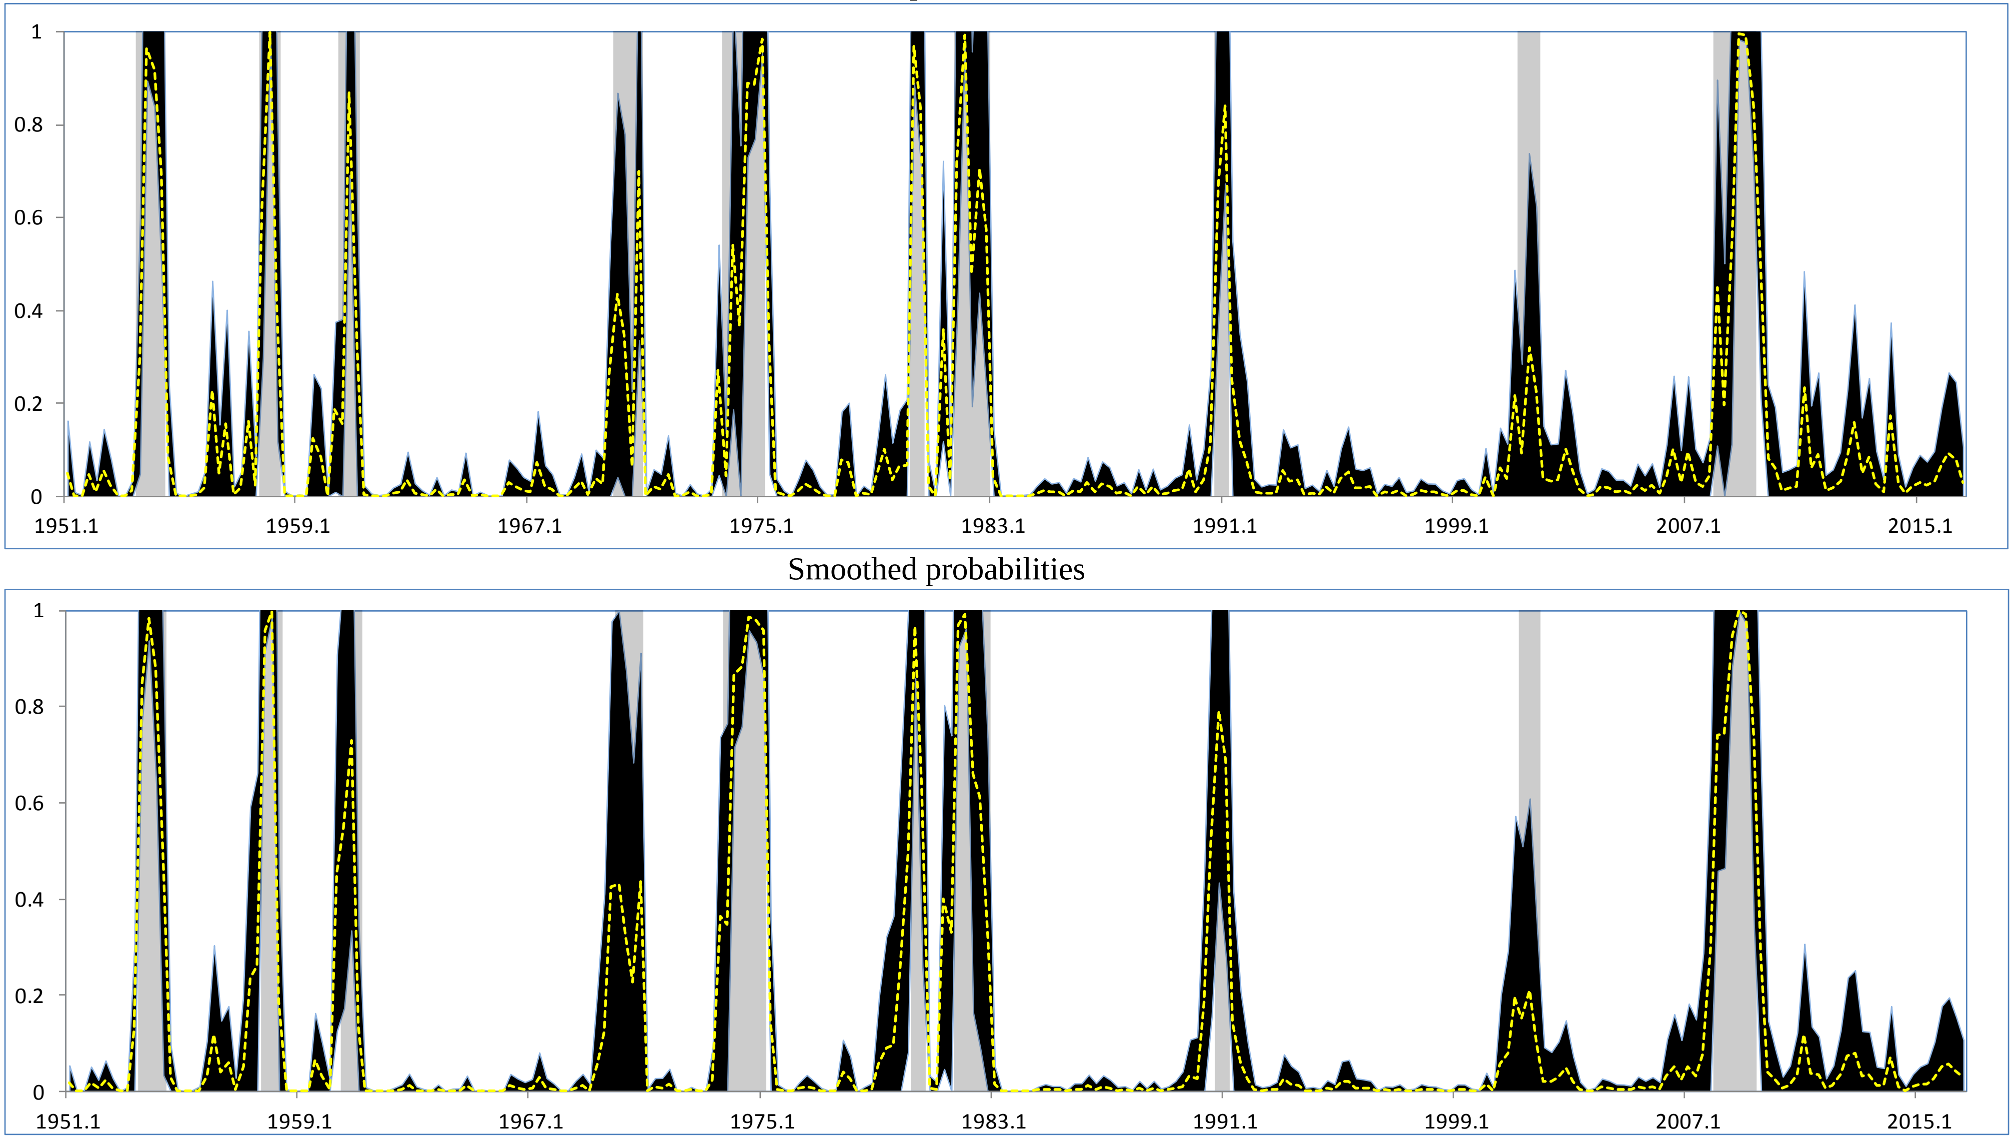

Notes. Dotted lines refer to the filtered (top panel) and smoothed (bottom panel) probabilities of recessions and black areas refer to their 95% confidence intervals. Grey areas refer to the NBER recessions.

Figure 7. Real-time filtered probabilities and confidence intervals

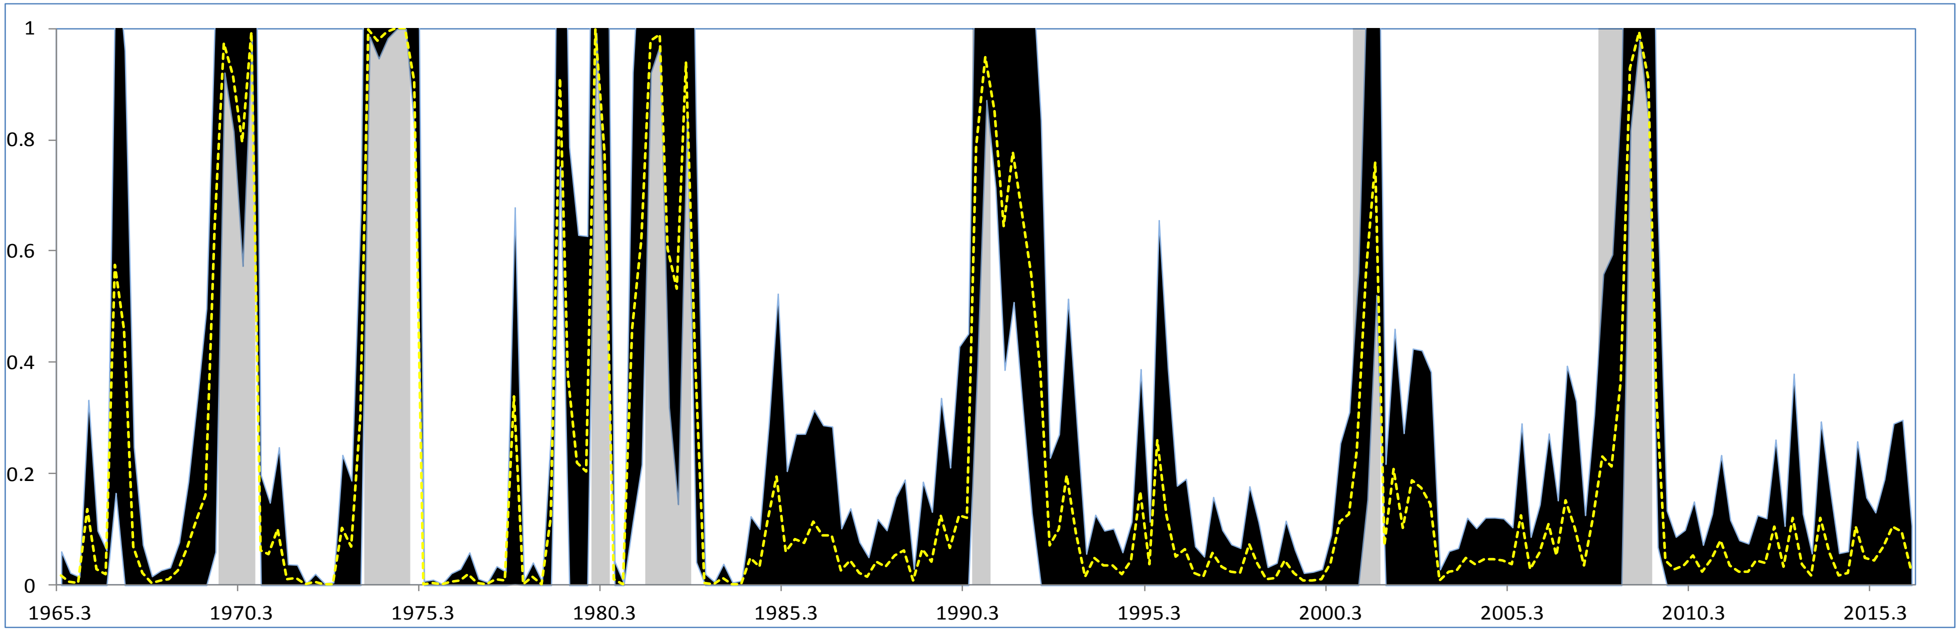

Notes. Dotted line refers to the real-time filtered probabilities of recessions and black areas refer to their 95% confidence intervals from 1965.3 to 2016.3. The Grey areas refer to the NBER recessions.

Figure 8. Ex-ante and ex-post real interest rate

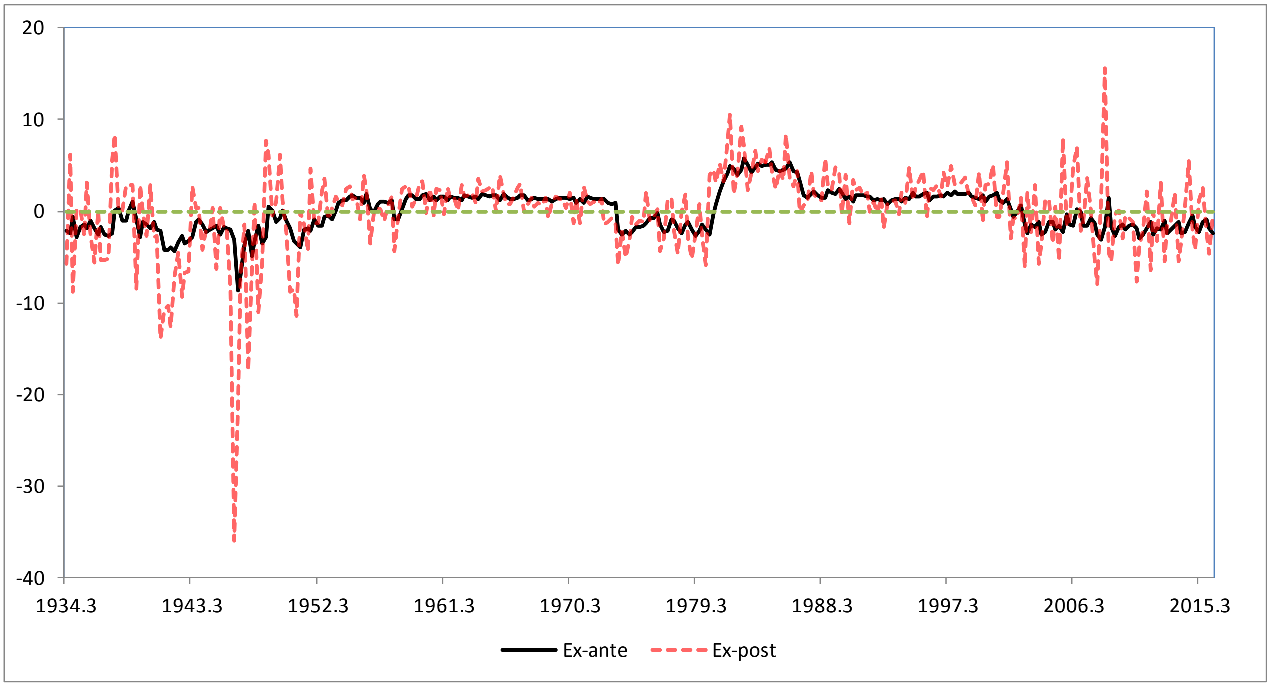

Notes. The ex-post interest rate is the difference between the end-of-quarter figures of nominal interest rate and the inflation rate, using data from 1934.1 to 2016.3. The ex-ante interest rate refers to its estimation in a three-state Markov-switching model with two lags where both means and variances are allowed to switch

Figure 9. Filtered probabilities and confidence intervals

Panel A. Regime with low interest rate

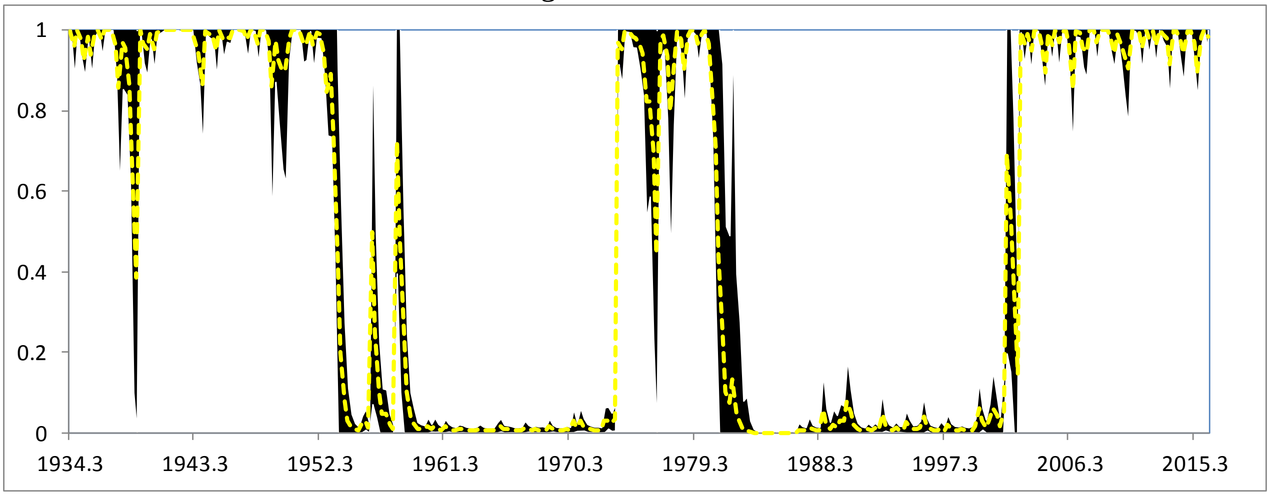

Panel B. Regime with medium interest rate

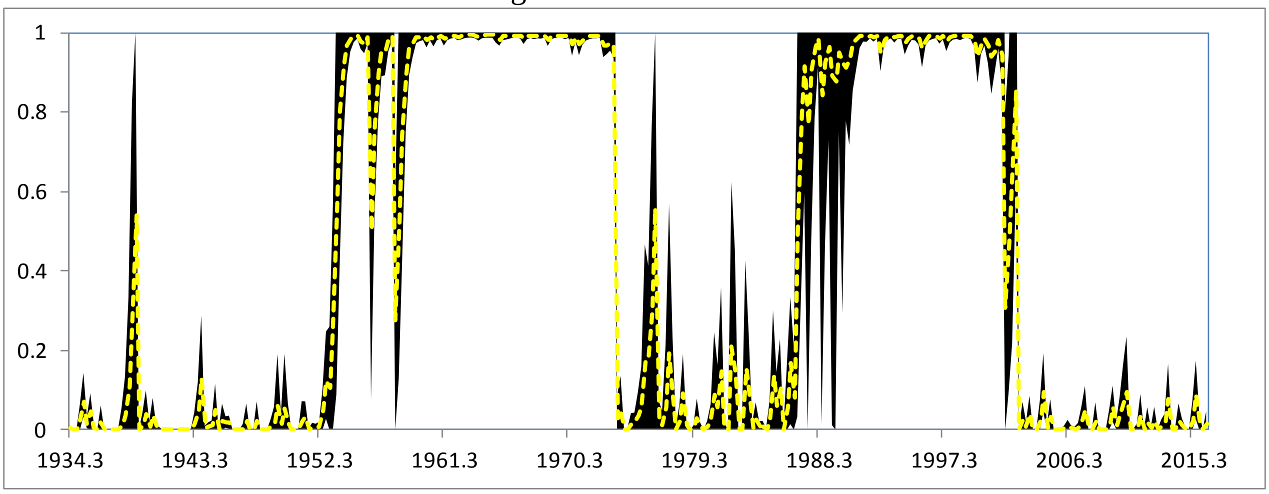

Panel C. Regime with large interest rate

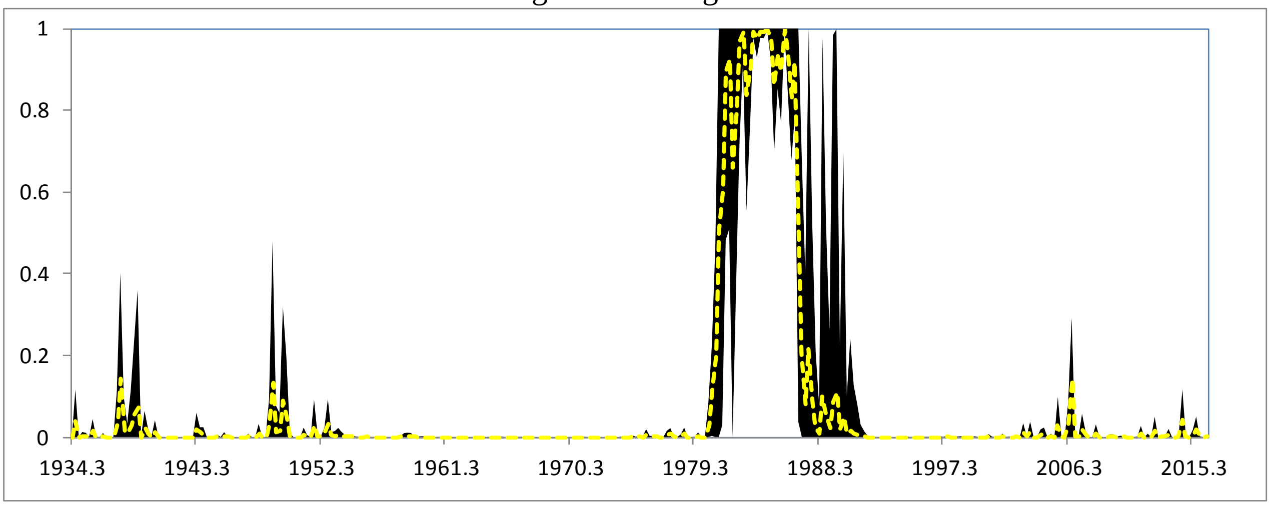

Notes. Dotted lines refer to the filtered probabilities of low (top panel), medium (middle panel), and large (bottom panel) interest rate. Black areas refer to their 95% confidence intervals.